

# ECSIC: Epipolar Cross Attention for Stereo Image Compression

Matthias Wödlinger  
TU Wien

Jan Kotera  
UTIA CAS,  
Czechia

Manuel Keulevic  
TU Wien

Jan Xu  
Deep Render

Robert Sablatnig  
TU Wien

## Abstract

In this paper, we present ECSIC, a novel learned method for stereo image compression. Our proposed method compresses the left and right images in a joint manner by exploiting the mutual information between the images of the stereo image pair using a novel stereo cross attention (SCA) module and two stereo context modules. The SCA module performs cross-attention restricted to the corresponding epipolar lines of the two images and processes them in parallel. The stereo context modules improve the entropy estimation of the second encoded image by using the first image as a context. We conduct an extensive ablation study demonstrating the effectiveness of the proposed modules and a comprehensive quantitative and qualitative comparison with existing methods. ECSIC achieves state-of-the-art performance in stereo image compression on the two popular stereo image datasets Cityscapes and InStereo2k while allowing for fast encoding and decoding.

## 1. Introduction

Stereo image compression (SIC) aims to jointly compress the two images of the stereo pair more efficiently by exploiting their mutual information and with the same goals as in lossy compression of single images – to reduce the bit rate required for storage and transmission while preserving the content and perceived quality. Stereo cameras are used in applications that demand high compression rates and low encoding and decoding latency to facilitate continuous recording or streaming (e.g., in autonomous driving or virtual reality streaming). As such, learned compression methods, which are typically symmetric in encoding and decoding time, offer an advantage over conventional methods, in which encoding is considerably slower. Learned methods transform the image to a latent representation which is then subsequently quantized and entropy coded to the bitstream using a learned probability distribution. They are trained end-to-end by optimizing the weighted trade-off between bitrate and distortion. Optimal rate is achieved by minimizing the cross-entropy of the true latent distribution and its

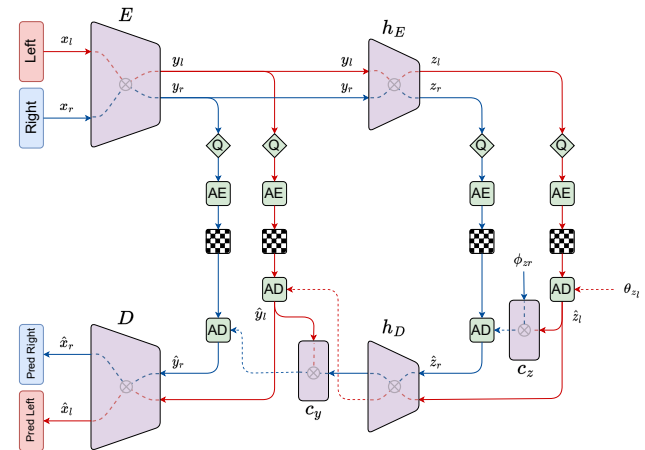


Figure 1. An overview of the architecture of our ECSIC model. The left and right streams are colored red and blue respectively. The encoder  $E$ , decoder  $D$ , hyperprior encoder  $h_E$ , and decoder  $h_D$  are jointly process the left and right image stream and run in parallel. The stereo context modules  $c_y$  and  $c_z$  are included only in the right stream and use input from the left stream as contextual information. The Stereo Cross Attention (SCA) modules connect the left and right image stream (depicted as  $\otimes$ ). Submodules in green (quantizers (Q) and arithmetic encoder/decoder (AE/AD)) do not contain any trainable parameters. The bitstreams are denoted with a checkerboard pattern. Dashed lines connecting to AD indicate predicted entropy parameters.

model.

State-of-the-art learned methods already outperform traditional methods, such as BPG, in single image compression [28]. Images in a stereo pair have high mutual information, optimal compression method should therefore achieve a rate close to that required to compress one of them and thus substantially outperform independent compression of the two images separately. However, occlusions and non-overlapping fields of view between the two cameras in a stereo setup make such a significant reduction of bitrate hard to achieve. To address these difficulties, early learned compression approaches focused on modelling the

disparity between the two images to obtain a dense warp field [24] or a rigid homography transform [11] to effectively register the two images in order to exploit their similarity. However, such an approach is quite computationally expensive, not to mention that dense disparity map cannot be efficiently transmitted via the bitstream. Later, more efficient approaches appeared that are able to achieve similar performance, such as SASIC [42], which combines a simple disparity model (global horizontal shift) with stereo attention to account for smaller local shifts.

In this work, we propose a novel neural network that compresses stereo images without explicit estimation of the disparity warping. The network follows an autoencoder structure with a hyperprior [4] entropy model. The encoder and decoder modules contain a novel stereo attention module that enables joint processing of both images in a stereo pair. Additionally, we introduce two stereo context modules in the entropy model for improved estimation by using the left image as context for the right image. The main contributions of this work are as follows:

- We propose ECSIC (Epipolar Cross attention for Stereo Image Compression), a novel stereo compression method that achieves state-of-the-art performance among SIC models while being fast during both encoding and decoding.
- We propose a stereo cross attention module and two stereo context modules for exploiting the mutual information in stereo images for compression and perform an ablation study demonstrating their impact on the overall stereo compression performance.
- We evaluate our method quantitatively and qualitatively on two popular stereo image datasets: Cityscapes and InStereo2k.
- The method is end-to-end trainable on any stereo image dataset and the code is publicly available <https://github.com/mwoedlinger/ecsic>.

## 2. Related Work

Based on their principle, image compression methods can be categorized into traditional and learned. In the former, the transformation from the input image to its latent representation is designed by hand; in the latter, it is learned from data by optimizing the rate-distortion loss. Common to both approaches is that the rate savings are ultimately achieved by using an (off-the-shelf) entropy coder that transforms the discrete latent representation to and from the minimum length bitstream. The approximate inverse of the initial transform is then used to reconstruct the decoded image.

**Traditional Methods** The best-known and most widely used image codec is arguably the JPEG method [41] from the 90s. It uses fixed 8x8 block tiling, chroma subsampling, discrete cosine transform, and several next-block prediction modes. Its successor JPEG2000 [34] is based on multi-resolution processing with a discrete wavelet transform. The development of modern compression methods has focused on video, and image codecs usually appear as wrappers around intra-frame compression in video codecs such as BPG [6] (based on HEVC [35]), AVIF [1] (based on AV1), or VVC-intra [8]. The latter has arguably the best compression performance among the traditional methods but is far from being adopted in practice due to its low speed, the lack of production-ready decoders, and restrictive licensing.

**Learned Methods** Pioneering work in learned image compression was done by Toderici et al. [38] who proposed a recurrent neural network for variable rate image compression. Foundations of the modern approach were laid by Ballé et al. [3] where an autoencoder-based model with a fixed parametrized distribution of the latent is trained with a rate-distortion loss for a fixed target bitrate. In their subsequent works, the fixed latent entropy model was replaced by a per-pixel Gaussian distribution with parameters predicted for each input image by a separate hyperprior module [4] or even an autoregressive context module [25, 28], which significantly improved performance.

This basic structure was later improved in many ways, for example, by changes in model architecture [9, 12, 17, 43], quantisation approximation [16, 37], or theoretical insights into the optimisation problem [44]. Much effort has been devoted to improving the context model [15, 17, 19, 23, 29, 32]. Recently proposed methods report performance improvements by replacing convolution blocks with transformers or other types of attention modules, mainly in the hyperprior and context model [22, 23, 31], but also in the main autoencoder [46, 47].

A slightly different line of research aims at the realism of image reconstruction in addition to the objective of rate-distortion optimisation. This is most commonly achieved by GANs [2, 13, 18, 26, 39] but also by other generative models such as denoising diffusion [14, 36].

**Stereo Image Compression** In compression of stereo images, bitrate is saved by exploiting the mutual information between the left and right images of the stereo pair. Although somewhat similar to compression of consecutive frames of a video sequence, disparity in stereo images is not well modelled by optical flow, and direct application of video codecs is therefore suboptimal. Of the traditional methods, MV-HEVC [27] is an extension of the HEVC video codec for multi-view sequences with good perfor-

mance but lacking support for high bit depth processing and 444 chroma mode. Learnable lossless stereo compression has been proposed by Huang et al. [21], based on explicit disparity estimation and image warping.

Several learned methods have been proposed for lossy stereo compression. Liu et al. [24] proposed the DSIC method, which uses a conditional entropy model where skip modules feed disparity warped features from the encoded first image into the second. Deng et al. [11] proposed the HESIC method, in which the second image is warped by an estimated homography, and only the residual is encoded. In addition, a context-based entropy model and a final quality enhancement module are used to reduce the bitrate and increase the reconstruction quality. This has been simplified in the SASIC method proposed by Wödlinger et al. [42], where the transformation between the images in the stereo pair is approximated by a channel-wise horizontal shift and only the residual for the second image is encoded. This is further enhanced by using stereo attention between the two images in the common decoder. Cross-attention in the decoder is also used in the recently proposed distributed multi-view method LDMIC by Zhang et al. [45]. Contrary to our method, they employ global encoder-to-decoder cross-attention and a single-image autoregressive entropy model. Mital et al. [30] propose a similar approach for the special case of distributed source coding where a correlated image is available during decoding.

Unlike HESIC or SASIC, the proposed method does not seek a parametric transformation between the stereo images and does not use residual coding. Different from LDMIC, our stereo attention module only attends to the epipolar line, resulting in a significant runtime reduction without any penalty in the performance. Another difference to LDMIC is that our method does not use autoregressive context and is therefore much faster in decoding. Instead, we condition the entropy model of the right image on the left image, effectively using the left image as a context without a significant runtime penalty.

### 3. Method

The proposed method follows the common structure [4] consisting of the main autoencoder and hyperprior, to which we add two non-autoregressive context modules; see overview in Fig. 1. In the main branch, consisting of the encoder  $E$  and decoder  $D$ , the input image pair  $(\mathbf{x}_l, \mathbf{x}_r)$  is transformed to the latent representation  $(\mathbf{y}_l, \mathbf{y}_r)$  and quantized to discrete tensors  $(\hat{\mathbf{y}}_l, \hat{\mathbf{y}}_r)$  which constitutes the bitstream. The decoder  $D$  reconstructs the output images  $(\hat{\mathbf{x}}_l, \hat{\mathbf{x}}_r)$ . In the hyperprior branch the hyper-encoder  $h_E$  transforms the latents to  $(\mathbf{z}_l, \mathbf{z}_r)$  which again undergo quantization to  $(\hat{\mathbf{z}}_l, \hat{\mathbf{z}}_r)$  and are stored in the bitstream as side-information. They are then used by the hyper-decoder  $h_D$  to estimate the entropy parameters of the

latents  $(\hat{\mathbf{y}}_l, \hat{\mathbf{y}}_r)$ . All these modules jointly process the left and right images in parallel.

We add two non-autoregressive *stereo context modules*  $c_y$  and  $c_z$  to the right image stream, which aid in estimating the entropy parameters of the right image latents  $\hat{\mathbf{y}}_r$  and hyperlatents  $\hat{\mathbf{z}}_r$ , respectively, using the information already available from the left side. The left and right image stream are further connected via the proposed *Stereo Cross Attention* (SCA) modules (see Section 3.2), which are included in all the modules that connect both streams – in encoder/decoder and hyper-encoder/hyper-decoder – and also in the stereo context  $c_y$ .

The resulting method can be trained end-to-end with the rate-distortion loss (see Sec. 3.4) on any dataset of stereo image pairs. Unlike other recent methods [11, 45], the proposed method does not include any autoregressive components, which allows for fast encoding and decoding (see Sec. 4.4).

### 3.1. Encoding/Decoding and Quantization

The encoder  $E$  and decoder  $D$  each consist of four convolutional layers with three down/up-sampling steps for left and right each, one SCA module, and PReLU [20] activation functions. The encoder and decoder of the hyper-prior similarly employ three convolutional layers with two down/up-sampling steps, one SCA module, and PReLU activation functions. The initial convolutions in the encoder  $E$  share weights between left and right streams. Network diagrams for each module can be found in the supplementary material. Each quantization operation applies integer rounding to the mean-subtracted input. For example for the left latent:

$$\hat{\mathbf{y}}_l = \text{round}(\mathbf{y}_l - \boldsymbol{\mu}_l) + \boldsymbol{\mu}_l, \quad (1)$$

where  $\boldsymbol{\mu}_l$  is the estimated mean of the distribution of  $\mathbf{y}_l$ . Analogously for  $\mathbf{y}_r, \mathbf{z}_l$ , and  $\mathbf{z}_r$ .

### 3.2. Stereo Cross Attention Module

We propose a new Stereo Cross Attention (SCA) module to facilitate the flow of non-local information between the left and right image compression stream. It performs cross attention between the corresponding epipolar lines – for each location in one image the attention domain is the corresponding horizontal row in the other image. By restricting the attention only to horizontal lines (under the assumption that the input images are rectified) we circumvent the issue of the quadratic memory complexity of vanilla attention and can process all rows in parallel. The resulting method still has quadratic complexity but only in the width rather than the total pixel count  $\mathcal{O}(w^2h)$ . The structure of the SCA module is shown in Fig. 2. Layer norm is applied only to queries and keys (not values, which constitute the final output). In the Multi-Head Attention (MHA) block, we

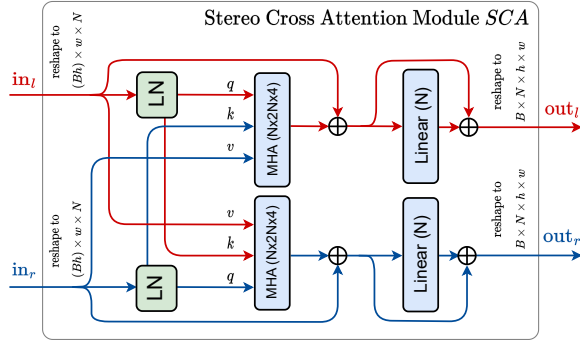


Figure 2. The proposed Stereo Cross Attention (SCA) module. The left and right streams are colored red and blue, respectively. LN denotes layer norm and MHA denotes multi-head attention block with arguments (output dimension  $\times$  embedding dimension  $\times$  heads). The streams denoted  $q$ ,  $k$ , and  $v$  refer to the standard query, key and value terminology.

use 1D convolutions with a kernel size of 3 instead of linear embeddings. We also tried different variants of positional encoding [33, 40] but found no impact on overall performance. The SCA module is included in all of the submodules that connect both streams –  $E$ ,  $D$ ,  $h_E$ ,  $h_D$ , and also in the stereo context  $c_y$ ; see Fig. 1. In  $E$  and  $h_E$ , the SCA module is applied after all downsampling layers and before the final convolutional layer. In  $D$  and  $h_D$ , the module is applied after the initial upsampling layer. Additional details can be found in the supplementary material.

### 3.3. Entropy model

Following the hyperprior structure from [4] we employ a pair of hyperlatents  $\hat{z}_l, \hat{z}_r$  as side information to aid in estimating the entropy parameters of the main latents  $\hat{y}_l, \hat{y}_r$ . In the following paragraphs, we write tensors (non-scalars) in bold, use  $\theta_{(\cdot)}$  for entropy parameters and  $\phi_{(\cdot)}$  for other learnable or predicted parameters.

The distribution of the left hyperlatent  $\hat{z}_l$  is modeled by a channel-wise Laplacian distribution  $\text{Lap}_{\mu, b}$  with parameters  $\theta_z^l := (\mu_z^l, b_z^l)$  for each channel of  $\hat{z}_l$  that are learned during training and fixed afterwards. The distribution of the right hyperlatent is modeled by a factorized Laplace distribution with parameters  $\theta_z^r := (\mu_z^r, b_z^r)$  for each pixel. These are predicted adaptively for each input. Likewise, the distribution of the main latents  $\hat{y}_{l/r}$  is also modeled by a factorized Laplace distributions with parameters  $(\theta_y^l, \theta_y^r)$ .

To reduce the bitrate we condition the right image entropy model on information from the left stream. To this end, we include two stereo context modules,  $c_y$  and  $c_z$ ; see Fig. 3.

The left hyperlatent entropy parameters are learned. The

right hyperlatent entropy parameters  $\theta_z^r$  are predicted by  $c_z$  from  $\hat{z}_l$  and a set of fixed (learnable) parameters  $\phi_{z_r}$ :

$$\theta_z^r = c_z(\hat{z}_l, \phi_{z_r}). \quad (2)$$

During encoding (decoding),  $\hat{z}_l$  is encoded (decoded) first using its fixed entropy model and then used to encode (decode)  $\hat{z}_r$  using entropy parameters predicted by  $c_z$ .

The parameters  $\theta_y^l$  of the distribution of the left latent  $\hat{y}_l$  are predicted from both hyperlatents  $\hat{z}_l, \hat{z}_r$  by the hyperprior decoder  $h_D$ . Similarly to the previous case, we use the decoded left latent  $\hat{y}_l$  to more accurately estimate the entropy parameters of the right latent. To this end we include the context module  $c_y$  which predicts the right latent entropy parameters

$$\theta_y^r = c_y(\hat{y}_l, \phi_{y_r}) \quad (3)$$

from the already decoded left latent and the second output  $\phi_{y_r}$  of the hyper-decoder  $h_D$ .

### 3.4. Loss Function

In training we optimize the rate-distortion loss

$$\mathcal{L} = \mathcal{R} + \lambda \mathcal{D}, \quad (4)$$

where  $\mathcal{R}$  denotes the rate and  $\mathcal{D}$  the distortion loss term;  $\lambda \in \mathbb{R}$  is a trade-off parameters that determines the average bitrate of the trained model. The distortion loss term is the expectation of the mean-squared-errors

$$\mathcal{D}(x_l, x_r) = \mathbb{E}_{x_l, x_r \sim p_w} [\|x_l - \hat{x}_l\|_2^2 + \|x_r - \hat{x}_r\|_2^2]. \quad (5)$$

To estimate the rate, we compute the cross entropy between the predicted distribution of our entropy model and the true distribution of the latents/hyperlatents. The total rate loss is then the sum of the rates of the latents and the hyperlatents:

$$\begin{aligned} \mathcal{R} = \mathbb{E}_{x_l, x_r \sim p_w} [ & -\log_2 p(\hat{z}_l | \theta_z^l) \\ & -\log_2 p(\hat{z}_r | \phi_{c_z}, \phi_{z_r}, \hat{z}_l) \\ & -\log_2 p(\hat{y}_l | \phi_{hd}, \hat{z}_r, \hat{z}_l) \\ & -\log_2 p(\hat{y}_r | \phi_{c_y}, \phi_{hd}, \hat{y}_l, \hat{z}_r, \hat{z}_l)], \end{aligned} \quad (6)$$

where  $\phi_{hd}, \phi_{c_y}, \phi_{c_z}$  denote the parameters of the hyperprior decoder and our proposed stereo context modules  $c_y$  and  $c_z$ , respectively, and  $p(\dots)$  are the Laplace distributions specified in the preceding section 3.3.

Since quantization has zero derivative almost everywhere it needs to be replaced by some proxy expression during training. As in [3] we approximate quantization with additive uniform noise for the rate loss (similarly for  $y_r, z_l$  and  $z_r$ )

$$\tilde{y}_l = y_l + \mathcal{U}(-0.5, 0.5). \quad (7)$$

Following Minnen et al. [29], we employ straight-through-estimation quantization during training for the distortion loss.

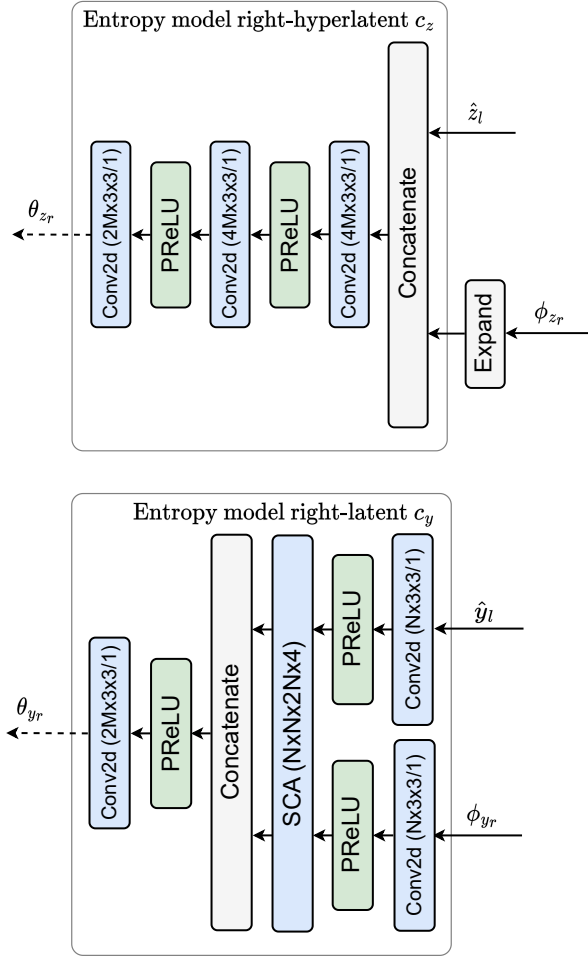


Figure 3. The context modules  $c_z$  and  $c_y$ . Top:  $c_z$  predicts the entropy parameters of the right hyperlatent  $\hat{z}_r$  from the left hyperlatent  $\hat{z}_l$  and a set of learned parameters  $\phi_{z_r}$ . Bottom:  $c_y$  predicts the entropy parameters of the right latent  $\hat{y}_r$  from the left latent  $\hat{y}_l$  and partial output of the hyper-decoder  $h_D$  denoted  $\phi_{y_r}$ . In our experiments, we set  $N = 192$  and  $M = 48$ . The arguments after convolutions denote (output dimension  $\times$  kernel / stride).

## 4. Experiments

We give a brief overview of the datasets on which we evaluate our method, followed by implementation details. We then show rate-distortion curves, an analysis of encoding/decoding times, and conclude the section with an ablation study.

### 4.1. Experimental Setup

**Datasets:** We evaluate our method on two popular stereo image datasets, Cityscapes [10] and InStereo2k [5]. The datasets capture two different settings for stereo imagery. Cityscapes contains stereo pairs of driving scenes with large

variations of disparities in a single image pair. The InStereo2k dataset, on the other hand, contains only indoor scenes of assortments of objects closer to the camera. The Cityscapes dataset contains 5000 stereo image pairs of urban street scenes taken while driving in German cities. The images have a resolution of  $2048 \times 1024$  and are divided into 2975 training, 500 validation, and 1525 test image pairs. Following conventions [42, 45] we crop 64, 256 and 128 pixels from the top, bottom and sides respectively to remove car parts and artefacts from the rectification process. The InStereo2k dataset contains 2060 stereo images of indoor scenes. The images have a resolution of  $1080 \times 860$  and are divided into 2010 training and 50 test image pairs. We crop the images symmetrically so that the height and width are multiples of 32.

**Implementation Details:** In all experiments we set the number of channels in the encoding and decoding modules to  $N = 192$  and the number of latent channels to  $M = 48$ . We use 4 headed attention blocks with an embedding dimension of 386. To generate the rate-distortion curves, we train our method for 450 epochs on cityscapes and for 650 epochs on InStereo2k ( $\sim 1.3M$  steps each). We use the Adam optimizer with an initial learning rate of  $10^{-4}$  for a batch size of 1 and reduce the learning rate to  $10^{-5}$  after 1M steps. We vary  $\lambda$  depending on the targeted bitrate. We train on random crops of size  $256 \times 1024$  and evaluate on full-resolution images (except for dataset-specific crops as specified above).

**Codecs:** We compare our method with an extensive list of conventional and learned codecs. The codecs can be broadly grouped into four categories: single-image compression (BPG), video compression (HEVC), multi-view compression (MV-HEVC, LDMIC), and stereo image compression (HESIC/HESIC+, DSIC, SASIC, ECSIC (ours)). BPG [6] is applied to each frame independently without chroma subsampling. The video codecs HEVC [35] and VVC [8] are applied to a stereo image pair as two-frame video sequence with chroma subsampling disabled as it would unnecessarily degrade the PSNR score. We use the official reference implementation<sup>1</sup> with main\_444\_12 profile (YUV444 12bit) for HEVC. For VVC we report the values from Zhang et al. [45] where the lowdelay\_p configuration as well as YUV444 input format is used. MV-HEVC<sup>2</sup> [27] is used in the two-view intra-mode configuration (unfortunately, it only supports 4:2:0 chroma mode, resulting in worse PSNR scores at higher bitrates). We also report scores for the learned stereo compression methods DSIC [24], HESIC+ [11], and SASIC [42] from the respective papers. For LDMIC [45] we show the reported scores

<sup>1</sup><https://vcgit.hhi.fraunhofer.de/jvet/HM>

<sup>2</sup><http://hevc.info/mvhevc>

for the full LDMIC method, which includes an autoregressive context model, and a smaller version LDMIC (fast) without the autoregressive components.

## 4.2. Results

The Peak Signal to Noise Ratio (PSNR) rate-distortion curves can be seen in Fig. 5. We also report the Bjøntegaard Delta bitrate (BD-Rate) and BD-PSNR scores [7] for each codec w.r.t. BPG for Cityscapes and InStereo2K in Tab. 1 and Tab. 2, respectively. On InStereo2k, the proposed method outperforms all other codecs tested. On Cityscapes, Our method performs worse than VVC for low bpp (bits per pixel) ( $< 0.15$ ) but is the first learned method to outperform VVC in PSNR for  $\text{bpp} > 0.15$ . All other codecs are outperformed by our method. The second best learned method LDMIC relies on autoregressive entropy modelling in the default version which renders it much slower than our method (see Sec. 4.4). The fast variant of LDMIC without autoregressive context performs much worse with BD-Rate w.r.t. to ECSIC of 24.35% for InStereo2k and 49.23% for Cityscapes.

Our method also outperforms the other learned SIC models DSIC, HESIC+ and SASIC. Both HESIC+ and SASIC rely on explicit warping to remove spatial redundancies. HESIC+ uses a homography warping and SASIC applies a channel-wise shift to the latents. Interestingly, SASIC also uses a variant of cross attention in its decoder. The performance gap between SASIC and the proposed method shows the effectiveness of the proposed SCA and stereo context modules.

We provide a qualitative comparison of our method against a selected list of codecs on an image from the InStereo2k dataset in Fig. 4. Additional qualitative comparisons for samples from both the Cityscapes and InStereo2k datasets as well as MS-SSIM rate distortion curves can be found in the supplementary material.

## 4.3. Ablation Study

To assess the impact of individual components/additions in our method compared to a single image compression baseline, we generate rate-distortion curves on Cityscapes and InStereo2k for various architectural modifications of our method. The resulting rate-distortion curves are shown in Fig. 6.

**Baseline:** To assess the impact of individual components on stereo compression performance, we construct a baseline method by stripping the ECSIC model of both context modules  $c_y$  and  $c_z$  (the entropy modelling for left and right is independent of each other) and removing each SCA module in the remaining architecture (including the corresponding activation functions of each SCA layer). The result is

Table 1. Relative quality difference (PSNR gain at the same bitrate; higher is better) and bitrate difference (bitrate gain for the same PSNR; lower is better) of the benchmarked methods on Cityscapes w.r.t. BPG. Best results in bold and second best underlined.

Method	BD-PSNR [dB]↑	BD-Rate [%]↓
ECSIC (proposed)	<u>2.86</u>	<u>-51.90</u>
LDMIC	2.07	-42.20
LDMIC (fast)	1.35	-29.66
SASIC	0.98	-22.40
DSIC	0.07	-3.35
VVC	<b>3.12</b>	<b>-56.24</b>
HEVC	1.14	-25.78
MV-HEVC	0.41	-10.07
BPG	0.0	-0.0

Table 2. Relative quality difference (PSNR gain at the same bitrate; higher is better) and bitrate difference (bitrate gain for the same PSNR; lower is better) of the benchmarked methods on InStereo2K w.r.t. BPG. Best results in bold and second best underlined.

Method	BD-PSNR [dB]↑	BD-Rate [%]↓
ECSIC (proposed)	<b>1.57</b>	<b>-42.08</b>
LDMIC	<u>1.26</u>	<u>-41.03</u>
LDMIC (fast)	0.87	-30.40
SASIC	0.38	-15.43
HESIC+	0.37	-14.90
VVC	0.86	-31.02
HEVC	0.45	-15.09
MV-HEVC	0.19	-4.96
BPG	0.0	-0.0

two separate models that independently compress the left and right images of a stereo image pair.

**ECSIC (proposed):** The proposed method with all additions as shown in Fig. 1. The largest gains over the baseline are at low bit rates. For example, for Cityscapes, the BD-rate limited to low PSNR (34–38dB) shows a bitrate saving of 37.0%, while in the high PSNR range (44–46dB) the difference is reduced to 19.0%. The maximum asymptotic theoretical bitrate saving that can be achieved is 50.0%, which corresponds to compressing a stereo image pair with the bitrate of a single image. In reality, the optimum is even lower due to occlusions and non-overlapping fields of view in the stereo pair.

**Only encoder SCA:** We extend the baseline method by adding a single SCA layer to the encoder  $E$ . The resulting



Figure 4. A qualitative comparison of our method against other codecs on an image from the InStereo2K test set. The zoomed-in region is upscaled with nearest neighbour upsampling.




Figure 5. Rate-distortion curves of our method and other codecs on Cityscapes (left) and InStereo2K (right) datasets measured by PSNR.

model shows no significant improvement over the baseline method. However, we have found that adding SCA modules in the encoder does improve performance if corresponding SCA modules are present in the decoders.

**Only decoder SCA:** We extend the baseline method by adding a single SCA layer in the decoder  $D$ . Contrary to adding SCA only in the encoder, adding a single SCA layer in the decoder already gives an improvement of 11.7% over the baseline method on Cityscapes.

**No context modules:** Removing the context modules  $c_y$  and  $c_z$  results in a 12.5% rate reduction compared to the full ECSIC model.

Our comparison shows that both, the proposed context modules  $c_y$  and  $c_z$  as well as the proposed SCA modules enable better compression of stereo images when compared to a single image compression baseline model that compresses both images independently. Furthermore, our experiments suggest that the SCA module works best in the decoding parts of the model. However, we found that SCA modules in the encoding part of the model in conjunction with SCA modules in the corresponding decoding

parts of the model leads to the best overall performance. We also conducted experiments with different variants of positional encoding [33, 40] but found no impact on compression performance.

#### 4.4. Coding Complexity

Fig. 7 depicts the average encoding and decoding times of our method against other methods on the InStereo2k dataset (i.e. the images are already rectified). The conventional methods BPG, HEVC and MV-HEVC were evaluated on an Intel Xeon Gold 6230R processor with a single core (times are taken from Zhang et al. times [45]). For LDMIC we show their reported encoding and decoding times [45] (measured on an NVIDIA RTX 3090 GPU). For SASIC and our proposed ECSIC model we measure encoding and decoding times on an NVIDIA RTX 3090 GPU. To make the comparison with conventional methods running on CPU fair, we also include the time needed to load the image and move it on the same GPU as the model in the encoding time. On our machine, this takes on average about 5ms which, for ECSIC, is slightly less than one third of the total reported encoding time. The proposed method shows low encoding and decoding times, beating all other methods in this benchmark.




Figure 6. Rate-distortion curves for our method with varying modifications for Cityscapes (left column) and InStereo2K (right column) measured by PSNR.

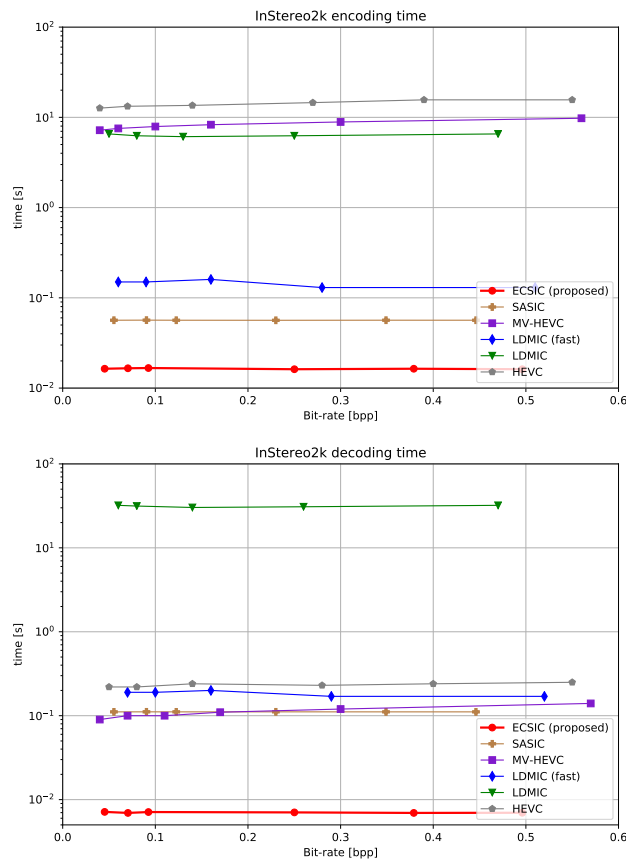


Figure 7. Encoding and decoding times of our method against other codecs on a logarithmic scale. The reported times are averages over the InStereo2k test set. All learned methods run on a GPU.

## 4.5. Limitations

Our evaluation is, as is typical in the field of stereo image compression, limited to particular available datasets where training and test sets are from the same distribution. It is expected that the resulting method will not be able to generalize to different camera setups if no additional care is taken during training (e.g., larger and more diverse training sets with varying camera setups and potentially concurrently increasing the model size). Furthermore, since our SCA module restricts the cross-attention operation essentially to the epipolar line, the input pairs must be rectified first. Finally, our method relies on a GPU or some other neural hardware accelerator to achieve the reported encoding/decoding times.

## 5. Conclusion

In this work, we have presented ECSIC, a novel method for stereo image compression. Our method consists of a convolutional neural network augmented with stereo attention modules that enable the network to compress both images jointly by exploiting the mutual information between them. Additionally, we proposed two stereo context modules for better entropy modelling of stereo images and showed in an ablation study the effectiveness of the proposed methods. The resulting model is fast in both encoding and decoding and outperforms all other learned compression methods on the two stereo image benchmark datasets Cityscapes and InStereo2k.

## 6. Acknowledgements

This project has received funding from the European Union’s Horizon 2020 research and innovation program under grant agreement No 965502.



## References

- [1] AVIF image format. <https://aomediacodec.github.io/av1-avif>, 2022. Accessed: 2023-03. 2
- [2] Eirikur Agustsson, Michael Tschannen, Fabian Mentzer, Radu Timofte, and Luc Van Gool. Generative adversarial networks for extreme learned image compression. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 221–231, 2019. 2
- [3] Johannes Ballé, Valero Laparra, and Eero P Simoncelli. End-to-end optimized image compression. In *5th International Conference on Learning Representations, ICLR 2017*, 2017. 2, 4
- [4] Johannes Ballé, David Minnen, Saurabh Singh, Sung Jin Hwang, and Nick Johnston. Variational image compression with a scale hyperprior. In *International Conference on Learning Representations*, 2018. 2, 3, 4
- [5] Wei Bao, Wei Wang, Yuhua Xu, Yulan Guo, Siyu Hong, and Xiaohu Zhang. Instereo2k: a large real dataset for stereo matching in indoor scenes. *Science China Information Sciences*, 63(11):1–11, 2020. 5
- [6] Fabrice Bellard. BPG Image format. <https://bellard.org/bpg>. Accessed: 2021-09-24. 2, 5
- [7] Gisle Bjontegaard. Calculation of average PSNR differences between RD-curves. *VCEG-M33*, 2001. 6
- [8] Benjamin Bross, Ye-Kui Wang, Yan Ye, Shan Liu, Jianle Chen, Gary J. Sullivan, and Jens-Rainer Ohm. Overview of the Versatile Video Coding (VVC) standard and its applications. *IEEE Transactions on Circuits and Systems for Video Technology*, pages 1–1, 2021. 2, 5
- [9] Zhengxue Cheng, Heming Sun, Masaru Takeuchi, and Jiro Katto. Learned image compression with discretized gaussian mixture likelihoods and attention modules. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7936–7945, 2020. 2
- [10] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016. 5
- [11] Xin Deng, Wenzhe Yang, Ren Yang, Mai Xu, Enpeng Liu, Qianhan Feng, and Radu Timofte. Deep homography for efficient stereo image compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1492–1501, June 2021. 2, 3, 5
- [12] Ge Gao, Pei You, Rong Pan, Shunyuan Han, Yuanyuan Zhang, Yuchao Dai, and Hojae Lee. Neural image compression via attentional multi-scale back projection and frequency decomposition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14677–14686, October 2021. 2
- [13] S. Gao, Y. Shi, T. Guo, Z. Qiu, Y. Ge, Z. Cui, Y. Feng, J. Wang, and B. Bai. Perceptual learned image compression with continuous rate adaptation. In *4th Challenge on Learned Image Compression*, Jun 2021. 2
- [14] Noor Fathima Ghouse, Jens Petersen, Auke Wiggers, Tianlin Xu, and Guillaume Sautière. Neural image compression with a diffusion-based decoder, 2023. 2
- [15] Zongyu Guo, Zhizheng Zhang, Runsen Feng, and Zhibo Chen. Causal contextual prediction for learned image compression. *IEEE Transactions on Circuits and Systems for Video Technology*, pages 1–1, 2021. 2
- [16] Zongyu Guo, Zhizheng Zhang, Runsen Feng, and Zhibo Chen. Soft then hard: Rethinking the quantization in neural image compression. In *International Conference on Machine Learning*, pages 3920–3929. PMLR, 2021. 2
- [17] Dailan He, Ziming Yang, Weikun Peng, Rui Ma, Hongwei Qin, and Yan Wang. ELIC: Efficient learned image compression with unevenly grouped space-channel contextual adaptive coding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5718–5727, June 2022. 2
- [18] Dailan He, Ziming Yang, Hongjiu Yu, Tongda Xu, Jixiang Luo, Yuan Chen, Chenjian Gao, Xinjie Shi, Hongwei Qin, and Yan Wang. PO-ELIC: Perception-oriented efficient learned image coding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 1764–1769, June 2022. 2
- [19] Dailan He, Yaoyan Zheng, Baocheng Sun, Yan Wang, and Hongwei Qin. Checkerboard context model for efficient learned image compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14771–14780, June 2021. 2
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015. 3
- [21] Zihao Huang, Zhe Sun, Feng Duan, Andrzej Cichocki, Peiying Ruan, and Chao Li. L3c-stereo: Lossless compression for stereo images. *arXiv preprint arXiv:2108.09422*, 2021. 3
- [22] Jun-Hyuk Kim, Byeongho Heo, and Jong-Seok Lee. Joint global and local hierarchical priors for learned image compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5992–6001, June 2022. 2
- [23] A. Burakhan Koyuncu, Han Gao, Atanas Boev, Georgii Gaikov, Elena Alshina, and Eckehard Steinbach. Contextformer: A transformer with spatio-channel attention for context modeling in learned image compression. In Shai Avidan, Gabriel Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner, editors, *Computer Vision – ECCV 2022*, pages 447–463, Cham, 2022. Springer Nature Switzerland. 2
- [24] Jerry Liu, Shenlong Wang, and R. Urtasun. DSIC: Deep stereo image compression. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3136–3145, 2019. 2, 3, 5
- [25] Fabian Mentzer, E. Agustsson, M. Tschannen, R. Timofte, and L. Gool. Conditional probability models for deep image compression. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4394–4402, 2018. 2

- [26] Fabian Mentzer, George D Toderici, Michael Tschannen, and Eirikur Agustsson. High-fidelity generative image compression. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 11913–11924. Curran Associates, Inc., 2020. [2](#)
- [27] Philipp Merkle, Karsten Muller, Aljoscha Smolic, and Thomas Wiegand. Efficient compression of multi-view video exploiting inter-view dependencies based on h. 264/mpeg4-avc. In *2006 IEEE International Conference on Multimedia and Expo*, pages 1717–1720. IEEE, 2006. [2](#), [5](#)
- [28] David Minnen, Johannes Ballé, and George D Toderici. Joint autoregressive and hierarchical priors for learned image compression. *Advances in Neural Information Processing Systems*, 31:10771–10780, 2018. [1](#), [2](#)
- [29] David Minnen and Saurabh Singh. Channel-wise autoregressive entropy models for learned image compression. In *2020 IEEE International Conference on Image Processing (ICIP)*, pages 3339–3343, 2020. [2](#), [4](#)
- [30] Nitish Mital, Ezgi Özyilkan, Ali Garjani, and Deniz Gündüz. Neural distributed image compression with cross-attention feature alignment. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2498–2507, 2023. [3](#)
- [31] Yichen Qian, Xiuyu Sun, Ming Lin, Zhiyu Tan, and Rong Jin. Entroformer: A transformer-based entropy model for learned image compression. In *International Conference on Learning Representations*, 2022. [2](#)
- [32] Yichen Qian, Zhiyu Tan, Xiuyu Sun, Ming Lin, Dongyang Li, Zhenhong Sun, Hao Li, and Rong Jin. Learning accurate entropy model with global reference for image compression, 2020. [2](#)
- [33] Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. Self-attention with relative position representations. In *Proceedings of NAACL-HLT*, pages 464–468, 2018. [4](#), [7](#)
- [34] A. Skodras, C. Christopoulos, and T. Ebrahimi. The JPEG 2000 still image compression standard. *IEEE Signal Processing Magazine*, 18(5):36–58, 2001. [2](#)
- [35] Gary J. Sullivan, Jens-Rainer Ohm, Woo-Jin Han, and Thomas Wiegand. Overview of the High Efficiency Video Coding (HEVC) standard. *IEEE Transactions on Circuits and Systems for Video Technology*, 22(12):1649–1668, 2012. [2](#), [5](#)
- [36] Lucas Theis, Tim Salimans, Matthew D. Hoffman, and Fabian Mentzer. Lossy compression with gaussian diffusion, 2022. [2](#)
- [37] Lucas Theis, Wenzhe Shi, Andrew Cunningham, and Ferenc Huszár. Lossy image compression with compressive autoencoders. In *International Conference on Learning Representations*, 2017. [2](#)
- [38] G. Toderici, Sean M. O’Malley, S. J. Hwang, Damien Vincent, David C. Minnen, S. Baluja, Michele Covell, and R. Sukthankar. Variable rate image compression with recurrent neural networks. *CoRR*, abs/1511.06085, 2016. [2](#)
- [39] Michael Tschannen, Eirikur Agustsson, and Mario Lucic. Deep generative models for distribution-preserving lossy compression. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. [2](#)
- [40] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. [4](#), [7](#)
- [41] G.K. Wallace. The JPEG still picture compression standard. *IEEE Transactions on Consumer Electronics*, 38(1):xviii–xxxiv, 1992. [2](#)
- [42] Matthias Wödlinger, Jan Kotera, Jan Xu, and Robert Szeliski. SASIC: Stereo image compression with latent shifts and stereo attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 661–670, 2022. [2](#), [3](#), [5](#)
- [43] Yueqi Xie, Ka Leong Cheng, and Qifeng Chen. Enhanced invertible encoding for learned image compression. In *Proceedings of the 29th ACM International Conference on Multimedia*, MM ’21, page 162–170, New York, NY, USA, 2021. Association for Computing Machinery. [2](#)
- [44] Yibo Yang, Robert Bamler, and Stephan Mandt. Improving inference for neural image compression. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 573–584. Curran Associates, Inc., 2020. [2](#)
- [45] Xinjie Zhang, Jiawei Shao, and Jun Zhang. LDMIC: Learning-based distributed multi-view image coding. *arXiv preprint arXiv:2301.09799*, 2023. [3](#), [5](#), [7](#)
- [46] Yin hao Zhu, Yang Yang, and Taco Cohen. Transformer-based transform coding. In *International Conference on Learning Representations*, 2022. [2](#)
- [47] Renjie Zou, Chunfeng Song, and Zhaoxiang Zhang. The devil is in the details: Window-based attention for image compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17492–17501, June 2022. [2](#)